

United States Application Entitled

A SYSTEM AND METHOD FOR HOST BASED STORAGE VIRTUALIZATION

Inventor: Thomas Jaskiewicz

Patent Application

A SYSTEM AND METHOD
FOR HOST BASED STORAGE VIRTUALIZATION

5 **Field of the Invention**

The illustrative embodiment of the present invention relates generally to network storage and more particularly to a host based storage method utilizing a virtual interface between a host electronic device and network storage mediums.

10

Background of the Invention

As the use of computers and computer networks has grown, so has the need for data storage. Duplicate copies of important data must be saved or "backed up" so that in the event of hardware or software failure the data is not lost. Additionally, large amounts of original data often must be saved in locations other than where the data is generated due to storage size considerations. Conventional methods of storing data have evolved from host based methods of storing data (i.e., copying data to multiple files and/or multiple disk drives on the same system) to network based storage solutions that store data at sites remote from the host.

20

A conventional network based storage system includes a volume controller and a multitude of storage mediums under the control of the volume controller. A volume is a storage medium partition which includes a table of contents listing all of the files stored in the partition. A storage medium may have more than one partition. A volume controller is a device used to send commands to the volumes in the storage system (i.e.: it enables access to the data in the volumes). The storage mediums are usually a collection of disk drives gathered in a single apparatus, such as a RAID system (RAID is an acronym for Redundant Array of Inexpensive/Independent Disks).

25

30

Most network storage systems are designed to attempt to provide continual access to data and continual protection of data. Unfortunately, the twin goals of protection and uninterrupted access sometimes conflict. In conventional network storage systems, the repair process to address hardware failure in the system interrupts access to data. Conventionally, a database or file system sends a collection of data over a network to a volume controller for storage. The volume controller sends the data to the storage mediums. This method of relying upon the volume controller to decide

35

where to physically store data requires the proper functioning of the storage elements of the system.

5 If a device holding a storage medium fails, the volume controller attempts to recover from the failure by locating and activating a spare drive ("spindle"). During the process of substituting a spare spindle for the failed storage medium, access to data is interrupted during the activating of the spindle and the copying of data to the spindle. If there is not a spare spindle that is locatable by the volume controller, the data in the storage system may be unprotected and/or inaccessible depending upon the type of storage system being employed. Likewise, in network storage systems which leave the determination of the ultimate storage destination of data to an element in the network (i.e.: software on a switch or other device remote from the host electronic device) the data is similarly inaccessible during network storage system repairs. Similar delays in access to data accompany the re-configuration of the storage mediums used in the network storage system. Users, such as a file system or database, which are directly interfaced with the storage system are required in conventional network storage systems to wait for data access during the reconfiguration and repair of the storage system.

Summary of the Invention

20 The illustrative embodiment of the present invention provides a method of inserting a virtual layer between the user of a host electronic device and a network storage medium. The virtual layer wraps a real volume of storage (i.e.: a logical unit) inside a virtual volume. The virtual logical unit (VLUN) is a software created virtual interface that encapsulates and hides some aspects of the real volume. One of the hidden attributes is the location of the real volume. The VLUN is implemented as a pseudo-device situated between the real volume and the user of the data in the volume. All accesses to the data in the real volume pass through the VLUN. By providing a virtual layer between the user and the network storage system, reconfiguration and repair of the storage system may be undertaken with minimal disruption of data access to the user.

The hiding of the real volume enables real-time data relocation. A user seeking access to the data through the VLUN is unaware of the actual location of the data. As long as the virtual interface (the VLUN) remains in an accessible location, the actual location of the data is irrelevant to the user. This allows the data to be moved around in

the storage system as required. Such movement of data is often required during hardware failure and while performing preventive maintenance on the storage system (i.e.: swapping drives in RAID sets). The illustrative embodiments of the present invention provide duplicate copies of data which remain accessible through the VLUN
5 even during the reconfiguration and repair of the RAID system and thus enable a user to access data at all times.

In one embodiment of the present invention, a host electronic device is interfaced with the network. A plurality of devices having storage mediums with device
10 controllers, such as a RAID system, are also interfaced with the network. A software created VLUN is located on the host electronic device and acts as a virtual interface between a user of the host electronic device and the storage mediums. The users of the host electronic device, such as a file system, or database, send data write and read requests to the VLUN. The VLUN sends data write and read requests to a VLUN
15 selected device controller for a device with a storage medium. The storage process beyond the VLUN is transparent to the user. The user has no direct interaction with the storage medium. In one implementation, the selected device is a RAID set, the device controller is a RAID volume controller, and the VLUN selects the RAID set based on a desired RAID level.

20
In another embodiment of the present invention, an apparatus interfaced with a network, includes a file system, processor and storage medium. The network is also interfaced with a plurality of storage devices which includes storage mediums. The apparatus further includes a software facility to create a virtual interface used to allocate
25 data to available storage mediums. The storage mediums are transparent to the user of the apparatus and all read and write requests for storage are made through the virtual interface created by the software facility.

Description of the Drawings

30
Figure 1 is a block diagram of an environment suitable for practicing an illustrated embodiment of the present invention;

Figure 2 is a block diagram of the hierarchical location of the virtual interface (VLUN);

Figure 3 is a flow chart depicting the sequence of steps followed by an illustrative embodiment of the present invention to store data using a RAID volume controller;

Figure 4 is a flow chart of an alternate embodiment of the present invention indicating the steps followed upon the detection of an error in one side of a RAID set.

Figure 5 is a flow chart depicting the sequence of steps followed by an embodiment of the present invention to store data with the the VLUN directly allocating the data to storage mediums; and

Figure 6 is a flow chart of the sequence of steps followed by the embodiment of **Figure 5** in recovering from a detected error.

Detailed Description

The illustrative embodiments of the present invention provide a method for performing host based storage virtualization. A software created virtual interface, known a virtual logical unit (VLUN), is inserted between a user, such as a file system or database, on a host electronic device, and the end storage mediums located, in most cases, remotely over a network. The insertion of the VLUN abstracts the process of storing data such that the end storage mediums locations are opaque to the user. The user writes data to the VLUN and reads data from the VLUN. The VLUN is located on the host electronic device interfaced with the network. The location of the VLUN on the host electronic device enables the repair and reconfiguration of the storage mediums without requiring user knowledge of the details of the process. The local control of the VLUN allows a user to have uninterrupted access to the user's data even during reconfiguration and repair of the storage mediums.

Figure 1 depicts an environment suitable for practicing the illustrative embodiment of the present invention. A plurality of hosts, host A 2, host B 6, and host C 7 are interfaced with a network switch 8. Host A 2 includes a file system 1, memory 3, a VLUN 4 located in the memory, and a processor 5. Also interfaced with the network switch 8 are a plurality of storage devices 10, 14, 18, and 22. Each storage device 10, 14, 18, and 22 includes a volume controller 11, 15, 19, and 23, as well as a storage medium such as disk drives 12, 13, 16, 17, 20, 21, 24, and 25. The volume controller is used to write and read data from the storage mediums. Those storage mediums may be entire disk drives, or multiple volumes within a single disk drive. The

file system 1 located on host A 2 sends data to the VLUN 4. The VLUN 4 redirects the data to an available storage device 10, 14, 18, and 22 selected by the VLUN. The illustrated embodiment of the present invention may use a RAID system for backing up data. For instance, the VLUN 4 may send a collection of data for storage via the
5 network switch 8 to the RAID volume controller 15 located on storage device 14. The RAID volume controller 15 may then mirror the data on a mirror set made up of disk drives 16 and 17. "Mirroring" is a technique used to distribute multiple copies of the same data to two or more disks. Any time the data is updated, the data is changed for both copies. Similarly, the VLUN 4 may be used to read stored data. The VLUN 4 may
10 receive a read request from the file system 1 and send the request to the RAID volume controller 15. The RAID volume controller 15 retrieves the data from the disk drives 16 and 17 and sends the data to the VLUN 4 which sends it to the file system 1. The VLUN 4 provides an interface between a host user and a storage medium so that the process of storing or reading the data on the medium is hidden from the user. The
15 VLUN 4 may be moved between hosts 2, 6 and 7.

Figure 2 depicts a hierarchy utilized by the illustrative embodiment of the present invention. A file system 30 sends write and read requests to the VLUN 32. The file system 30 is a system that the operating system of the host 2 uses to organize files
20 and directories. The VLUN 32 is a software created virtual interface between the file system 30 and the rest of the network storage system. The specific locations of the storage devices (the logical units), which may be locally or remotely located, are opaque to the file system which reads and writes data through the VLUN 32. The VLUN 32 is capable of opaquely mapping data to many different types of storage
25 devices in response to one request from the file system 30. The VLUN 32 sends its own write and read requests to a RAID volume controller 34 selected by the VLUN. The RAID volume controller 34 writes and reads data on available disks 36 and 38 in response to the requests received from the VLUN. Each entity in the hierarchical diagram is able to give commands only to the entity directly underneath in the
30 heirarchical diagram. Thus the file system 30 may not write data directly to the RAID volume controller 34. Since the specific RAID volume controller location is immaterial to the file system 30 (which accesses the RAID volume controller 34 through the VLUN 32), the VLUN may substitute a different RAID volume controller (not shown) for the RAID volume controller 34 when storing data. Similarly, the disks 36 and 38
35 respond to a user's request for data by sending the data to the RAID volume controller 34 rather than directly to the VLUN 32 or file system 30. Those skilled in the art will

recognize that there are many possible hierarchies within the scope of the present invention. For example, in systems without RAID volume controllers, the VLUN 32 may write and read data directly to the controllers for the disk drives 36 and 38. Similarly the requests to the VLUN 32 may originate with a database instead of a file system 30.

The VLUN 4 is responsible for determining an available storage location for data sent from a user, such as a file system 1 or database, of the host electronic device 2.

Figure 3 depicts the sequence of steps utilized by an illustrative embodiment of the present invention to store and then read data using a VLUN 4. A user on a host electronic device 2 sends data to the VLUN 4 for storage (step 42). The VLUN 4 determines available storage locations and then sends the data to a RAID volume controller (step 44) for an available RAID set. The RAID volume controller sends the data to a RAID set (step 46). Subsequently the VLUN 4 receives a request for data access from the user (step 48). The VLUN 4 retrieves the data from the RAID volume controller and forwards it to the requesting user (step 50). Those skilled in the art will recognize that alternate implementations of the present invention are possible. For example, the VLUN 4 may send the data to more than one RAID volume controller for storage, thus implementing different RAID levels for the same data.

The VLUN 4 acts to ensure continued data access for a user in the event operations to reconfigure or repair the storage mediums are required or advisable.

Figure 4 depicts the sequence of events followed by an illustrative embodiment of the present invention upon the detection of an error in one side of a RAID set. A user on a host electronic device 2 sends data to the VLUN 4 for storage (step 54). The VLUN 4 sends the data to a selected RAID volume controller (step 56). The RAID volume controller sends the data to a RAID mirror set where multiple copies of the data are stored (step 58). Depending upon the RAID level of the RAID set, the copies may be two complete copies of the data, or may be a complete copy and a parity copy. The copies may or may not be striped across the storage mediums. In some embodiments, two disk drives may be used for the RAID mirror set. In other embodiments three or more whole disk drives may be used to create a RAID set. Alternatively, segments of disk drives with adjacent logical memory addresses may be utilized by a RAID volume controller to create a software RAID set. Those skilled in the art will recognize that there are a multitude of possible components, and many different configurations of

components, that may be utilized without departing from the scope of the current invention.

As an illustrative example, subsequently to storing the data, the RAID controller may detect an error in one spindle of the RAID set. The VLUN 4 is notified of the error in the RAID set (step 60). Once the error in the spindle of the RAID set has been detected, the method of the present invention provides multiple ways to recover from the detected error. In one embodiment, the RAID set is repaired with a spare spindle (i.e., spare disk drive) (step 62). If the error is detected during an access request from the VLUN 4, a valid copy of the data is sent to the VLUN before or during the repairs to the spindle with the error. When the new spindle is in place, the data from the valid spindle of the RAID set is copied to the new spindle. In another embodiment of the present invention, data is retrieved from the valid spindle of the RAID set via the RAID volume controller and sent to the VLUN 4. The VLUN 4 verifies available storage locations and sends the data to a new RAID volume controller to be placed into a new RAID set (step 64). In another embodiment, data is retrieved by the VLUN 4 from the valid spindle of the RAID set using the RAID volume controller. Data is thereafter sent to a controller for a software RAID set which has been created using multiple storage mediums (step 66). Similarly, the VLUN 4 may schedule the movement of data from one RAID set to another so as to balance hardware use. Requests for data received during scheduled reconfiguration are treated as in the error detection process described above. Those skilled in the art will recognize that there are multiple error recovery and reconfiguration procedures possible within the scope of the present invention and those listed here are illustrative but not exhaustive.

25

In an alternative embodiment, the VLUN 4 sends data directly to available storage mediums. **Figure 5** depicts the sequence of steps followed to allocate data directly from the VLUN 4. The VLUN 4 receives data from a user for storage (step 70). The VLUN 4 checks the available storage mediums, such as disk drives, both local and remote (step 72) for available storage space. The VLUN 4 sends the data directly to the controllers for the storage mediums (step 74). There are a number of ways of storing the data within the scope of the current invention. The data may be striped across multiple disks (step 76) or the data may be mirrored among several disks (step 78). Those skilled in the art will recognize that the data may be striped and mirrored across multiple disks. Alternatively, complete duplicate copies of data or

35

parity data may be stored. In this embodiment, the VLUN 4 allocates the data directly instead of sending it to an intermediate controller, such as a RAID volume controller.

If the VLUN 4 is allocating data directly to the storage mediums without using a RAID volume controller, the VLUN is responsible for performing reconfiguration and repair operations on the storage mediums in a manner that minimizes data access disruptions to a user of the host electronic device 2. **Figure 6** depicts a sequence of steps followed by the VLUN 4 in an illustrative embodiment of the present invention upon detection of an error in a storage medium holding data directly allocated by the VLUN. The VLUN 4 allocates data directly to the storage medium in the manner discussed above (step 84). The VLUN then receives a user request to either read or write the data and attempts access to the stored data. Upon attempting access, the VLUN 4 detects an error in part of the stored data (step 86). The error may be mechanical error in the storage medium preventing the writing or reading of data, or the error may be a data error detected by the VLUN in comparing multiple copies of the data retrieved from the storage mediums which interferes with the reading of data. The VLUN 4 retrieves a valid copy of the data and sends it to the user (step 88). The VLUN 4 attempts the repair/reconfiguration of the detected error contemporaneously with retrieving the valid copy of the data or immediately thereafter. If the VLUN 4 is able to identify an additional available storage medium, the valid data may be copied to the identified storage medium (step 90). If the VLUN 4 is unable to identify a single available storage medium suitable for the valid data, the VLUN may segment the data and distributes the segments to multiple disks (step 92).

The illustrative embodiments of the present invention provide a virtual interface for the reading and writing of storage requests of a user of a host electronic device. By making the storage process transparent to the user, real-time reconfiguration and repair of storage mediums can take place without unduly disrupting user access to data. By providing the virtual interface on the host electronic device rather than locating the virtual interface out on the network, the user maintains greater access to data and more flexibility in recovering from hardware and software errors.

It will thus be seen that the invention attains the objects made apparent from the preceding description. Since certain changes may be made without departing from the scope of the present invention, it is intended that all matter contained in the above description or shown in the accompanying drawings be interpreted as illustrative and not

- in a literal sense. Practitioners of the art will realize that the system configurations depicted and described herein are examples of multiple possible system configurations that fall within the scope of the current invention. Likewise, the types of storage structures noted in the drawings and description are examples and not the exclusive
- 5 types of storage structures which may be employed within the scope of the present invention.